

Applicant: Elad et al.  
Serial No.: 09/919,702  
October 26, 2004  
Page 2

**Specification Amendments**

Page 1, Line 21.

**BACKGROUND OF THE INVENTION**

Users of computer-mediated resources always have particular goals when accessing those resources. The goals may be sharp (learn address of company) or fuzzy (be entertained) may be temporary (find a restaurant) or persistent (achieve and maintain financial independence), and may consist of multiple related or independent sub-goals. Constructing accurate models of a user's goals is a critical prerequisite to providing intelligent interaction with that user. Unfortunately, there is no monolithic, domain-independent body of knowledge that can accurately supply enough information concerning likely user mental states, to make a universal interface practical. In fact, every new capability that becomes available modifies the set of potential goals, plans, and tasks that is relevant to discourse. Consequently, a static set of models can never be satisfactory for long. User goals with respect to a given domain are tightly related to tasks that may be accomplished in that domain and to the referents or objects of those tasks. Thus, an ideal system would utilize domain-specific (or sub-domain)-specific (or sub-domain-specific) information to infer the user's mental state from his interaction, and would support easy addition of such information to an existing interface. Additionally, to be helpful, a user interface must consider the history of interaction, including previous user signals, goals and their outcomes, and must consider the information that was recently disclosed to the user, and the various ways of referring to that information. While the invention is applicable to all forms of human/computer communication, the main theoretical underpinnings are to be found in verbal discourse phenomena. Most of the following description refers to verbal discourse, but the invention contemplates applicability to virtually all non-verbal discourse as well, including mouse-actions, gestures, winks, etc. Similarly, system outputs are shown as text, tables, and graphs, but can also include generated speech, audible tones, blinking lights, and arbitrary transducers that stimulate sensory organs of the user.

Applicant: Elad et al.  
Serial No.: 09/919,702  
October 26, 2004  
Page 3

Page 7, Line 8,

4. A method dynamically maintaining a map of domain and domain transition adjacency, and using that map as an aid to plan-recognition and focus-recognition.

Page 8, Line 1,

12. A method for invoking several parse-related tasks concurrently, in order to examine their feasibility and results, in order to derive the likely meaning of a sentence user input.

Page 9, Line 3,

Figure 2 shows the User Interface Agency, the User Interface Coordination Agency, and adjacent modules of the present invention.

Page 10, Line 3,

Figure 16 shows an example of Simplified Strength/Necessity Belief Calculus.

Figure 17 shows an example of Bayesian Network Belief Calculus.

Figure 18 shows a scripting example.

Page 20, Line 3,

In a preferred embodiment, the domain cortex consists of a two-layer representation, "from" and "to", of commonly associated states and parameters. Not only are related concepts clustered, but related transitions among states are clustered, allowing common backbone sequences to emerge in the representation. It should be noted that the domain cortex also provides a mechanism for responding to ill-formed input. The system can attempt to match any well-formed *components* of such input against the domain-cortex clusters, and assume the meaning that most nearly matches the ill-formed input. Such repairs will always be noted in system interaction, so that the user does not assume that ill-formed input is, in fact completely acceptable. Arrow 502 transmits the top-ranked interpretation of the input to the Domain Planner 208. Arrow 514 shows the transfer of the semantic representation of user signals to the Generator module, where it would be is paraphrased. Arrow 503 shows tasks issued by the

Applicant: Elad et al.  
Serial No.: 09/919,702  
October 26, 2004  
Page 4

domain planner communicated to the Sub-domain Agent Manager 208. Results, in the form of axioms or non-textual objects, are communicated back through the system, starting with arrow 504. These results reach the Problem Solving Planner 210, which communicates them to the generator module 206.

Page 26, Line 20.

The user has an opportunity to select this assumption, which indicates a slight dislike for EDS, and to correct it if it is in error. He may be removing EDS from the portfolio merely because he is interested in seeing how much EDS contributes to the resulting financial outcomes, perhaps even to make a case *for* the inclusion of EDS. This feature of visible thought renders the invention potentially much more useful than interaction with a human expert, as the system provides greater default transparency into its rationale than is available with a human respondent.

Page 28, Line 4.

Figure 16 shows 1603, a rule that computes our belief that a substance is coffee, given 4 tests. Each test is characterized by strength (how much it reduces the remaining uncertainty, if the test is known to be completely true) and a necessity (how much it reduces the certainty, if the is known to be completely false). The rationale of the example goes something like this: a substance is somewhat more likely to be coffee if we find it in a mug, or if it is a hot liquid, but cold liquids in non-mug containers can also be coffee. On the other hand, if the liquid is not brown, it is very likely not coffee, and if it is tea (not(not tea) [ (?X is not tea) is false]) then we are very sure it is not coffee, thus the final two clauses have large necessity weights. In example A., we compute B<sub>4</sub>, given that we know all of the premises to be completely true. In example B., we compute the strength of belief in the conclusion, given that we are 100% sure of each of each of the premises, except for the third premise, which we are only 50% sure of. Note that if we were 0% confident that the substance was not tea, then we would be 0% confident that is coffee. Figure 17 shows a fragment of a bayesian belief network. This approach to computing likelihood is appropriate when some mixtures of base and conditional probabilities are available. Additionally, these probabilities are can be learned, from data, over a period of time. The

Applicant: Elad et al.  
Serial No.: 09/919,702  
October 26, 2004  
Page 5

example shows likelihoods for "planning for retirement", given that a user has asked about mutual funds.

Page 28, Line 32.

One promising area of current research is in probabilistic LTAGS, which allow probabilities to be introduced directly into the LTAG system. These might be considered as an alternative to the two-pass (parse and disambiguate) model that we use in the preferred embodiment, this would require a dynamic update of tag probabilities to match the discourse context. As an alternative to the LTAG approach, Head-driven Phrase Structure Grammar (HPSGs) might provide good results for some languages as would approaches based on Augmented Transition Networks (ATNs). These grammar formalisms can be (semi-) automatically translated, so it is possible to move much of the grammar from one form to another with little effort. The XTAG system employs a parser based on Earley's algorithm, but other parsing strategies might be used, including: chart parsers, recursive-descent parsers, bottom-up parsers, head-corner parsers, as well as Left-to-right scanning, Leftmost derivation (LL) parsers, Left-to-Right scanning with Rightmost derivation (LR) parsers, and Look-Ahead, Left-to-Right scanning, Rightmost derivation (LALR) parsers. Other, pure learning based, pattern-matching systems might also, eventually, offer reasonably good grammar building and parsing mechanisms. For instance, inductive logic programming has been used to learn a natural language grammar from a corpus of example language. In the area of multi-agent systems platforms, many technology alternatives are available, including FIPA-OS, and Zeus. As a substitute for our knowledge representation system, many existing approaches might be employed, including various modal logic and truth maintenance systems. Rather than using KMQL or FIPA style messages, XML representation and DOM<sup>i</sup> protocol might eventually offer greater interoperability. In any case, translations among the various messaging systems are being developed, eventually making these choices less restrictive. Rather than a bi-layer model of the domain cortex, an N-layer model would permit greater specificity of sequences. This extension might be useful in domains where long recurring sequences of operations are the common. There are many potential alternatives to the belief calculus methods sketched in Fig. 16 and

**Applicant: Elad et al.**  
**Serial No.: 09/919,702**  
**October 26, 2004**  
**Page 6**

Fig. 17, including, but not limited to: Dempster-Schafer evidence calculations; MYCIN-style certainty factors; Dragoni-Giorgini belief revision functions; Fuzzy-Logic inference; and Bonissone's bounded uncertainty inference method. These methods vary greatly in their assumptions, and requirements, and they can easily be used in concert, as different approaches are required to simulate human reasoning, and to accommodate differing levels of a priori probability estimates.